Credibility and Dynamics of Collective Attention

Tanushree Mitra
tmitra3@gatech.edu

Graham Wright
gwrong@gatech.edu

Eric Gilbert
gilbert@cc.gatech.edu

School of Interactive Computing & GVU Center
Georgia Institute of Technology
Atlanta, GA, USA

arXiv:1612.08440v1 [cs.SI] 26 Dec 2016

ABSTRACT

Today, social media provide the means by which billions of people experience news and events happening around the world. However, the absence of traditional journalistic gatekeeping allows information to flow unencumbered through these platforms, often raising questions of veracity and credibility of the reported information. Here we ask: How do the dynamics of collective attention directed toward an event reported on social media vary with its perceived credibility? By examining the first large-scale, systematically tracked credibility database of public Twitter messages (47M messages corresponding to 1,138 real-world events over a span of three months), we established a relationship between the temporal dynamics of events reported on social media and their associated level of credibility judgments. Representing collective attention by the aggregate temporal signatures of an event reportage, we found that the amount of continued attention focused on an event provides information about its associated levels of perceived credibility. Events exhibiting sustained, intermittent bursts of attention were found to be associated with lower levels of perceived credibility. In other words, as more people showed interest during moments of transient collective attention, the associated uncertainty surrounding these events also increased.

CCS Concepts

•**Human-centered computing** → **Empirical studies in collaborative and social computing;** *Social media;* Credibility; Twitter;

1. INTRODUCTION

Online social networks act as information conduits for real-world news and events [7], largely driven by collective attention from multiple social media actors [50]. Collective human attention drives various social, economic and technological phenomenon, such as herding behavior in financial markets [44], formation of trends [2], popularity of news [50], web pages [39], and music [41], propagation of memes [24], ideas, opinions and topics [40], person-to-person word-of-mouth advertising and viral marketing [24], and diffusion of product and innovation [4]. Moreover, it is the key phenomenon underlying social media reporting of emerging topics and breaking news [31]. However, unlike traditional news media—where information is curated by experienced journalists—social media news is unfiltered and therefore not subject to the same verification process as information presented by way of conventional sources. This naturally calls into question its credibility and the means with which to assess its credibility. Although scholars have increasingly expressed concern over the threats posed by digital misinformation, ranging from panic and violence incitement in society to libel and defamation of individuals or organizations [20], questions concerning the relationship between collective attention and information credibility have not been systematically quantified.

A fundamental attribute underlying any collective human behavior is how that behavior unfolds over time [3, 10]. Is there a relationship between allocation of collective attention and perceived credibility of events reported through social media? Do occasional bursts in collective attention—as more eyes and voices are drawn to the event's reportage—correspond to less certain information concerning the event? Uncovering the relationship between collective human behavior and information credibility is important for assessing the veracity of event reportage as it unfolds on social media. This relationship, if it exists, can provide insights into ways to disambiguate misinformation from accurate news stories in social networks—a medium central to the way we consume information [7] and one where digital misinformation is pervasive [11].

Empirical attempts at answering these questions in naturalistic settings have been constrained by difficulties in tracking social media posts in conjunction with judgments concerning the accuracy of the underlying information. Previous studies have instead focused on individual case studies involving specific news events [29, 26, 1], or have retrospectively studied a set of multiple prominent events [13, 11] which were known to contain misinformation. While useful, these approaches raise sampling concerns. In particular, they are based on the post-hoc investigation of events with known credibility levels, and thus select on the dependent variable [47]. Although these studies suggest the possibility of spikes in collective attention when false rumors propagate through social networks, the relation between collective attention and information credibility has not been systematically tested.

We tested this relation by analyzing data from the first large-scale, longitudinal credibility corpus, called CREDBANK [32]. The massive dataset was constructed by iteratively tracking millions of public Twitter posts. Twitter is a microblogging site where people write short time-stamped messages publishing their daily social activities [36] or discussing world events [22]. Hence, tweeting activity comprises regular circadian rhythms [16] intertwined with irregular bursts of activities corresponding to real-world news events [31]. In recent years, Twitter has become an attractive source for disseminating information pertaining to news events [22]. CRED-

BANK's credibility corpus is based on tracking all large-scale, real-world events surfacing on Twitter between October 2014 and February 2015, followed by credibility assessment through a verified human annotation framework. This iterative framework resulted in an experimental setup that captured accurate human judgments of credibility of social media information soon after it gained collective attention. It is important to note that while this process cannot arrive at the *truth* of the event reporting (perhaps an impossible epistemological task), it does capture expert-level human judgment at the time the event unfolded. CREDBANK contains 1,377 social media events collected over a period of three months, along with 66 million tweets nested within the event topics. The uniqueness of the dataset is evident not only from the systematic collection process but also from the range of the collected events. It contains, for example, objections to red cards thrown soccer matches, as well as the emergence of Ebola in West Africa.

Although the nature of this data limits causal inference, we were able to test the correspondence between collective attention and the level of information credibility. After filtering out unique event instances, we were left with a pruned corpus of 1,138 real-world events spread over 47M tweets. Analyzing this massive dataset, we find that the amount of recurring collective attention bursts could be used to determine the level of perceived credibility of an event. Specifically, we demonstrate that multiple occasional bursts of collective attention toward an event is associated with lower levels of perceived credibility. This finding opens a new perspective in the understanding of human collective attention and its relation to the certainty of information. In doing so, our results can have widespread implications in fields where predictive inference based on online collective interests dictates economic decisions, emergency responses, resource allocation or product recommendations [15, 51]; hence trusting the credibility of the collective reports is essential for an accurate anticipation by the predictive process.

2. RELATED WORK

2.1 Social Media and Credibility

With social media's rise as an important news source [7], individuals are constantly relying on online social networks to consume and share information, without recourse to official sources. However, modern online social networks like Facebook and Twitter are neutral with respect to the quality of information [13]. Moreover, users of these sites have been found to be poor judges of information credibility based on content alone [35]. Thus, scholars have increasingly become interested in assessing the credibility of social media content. Studies have focused on investigating specific events that were subjects of misinformation, such as the spread of rumors during the 2011 Great East Japan earthquake [26], the 2013 Boston marathon bombings [29] and the 2014 Sydney siege event [1]. Studies have also engaged in extensive analysis of multiple historically reported cases of rumor, such as, automatically classifying rumor instances [53] or predicting the credibility level of tweets [6, 18, 38]. However, these studies are based on the retrospective investigation of popular historical events. Hence, they suffer from the *selection on the dependent variable* confound. On the contrary, our study overcomes this confound by grounding its results on CREDBANK's data [32] – a credibility corpus which asks human raters to assess the credibility of *all* social media events in near real-time.

2.2 Collective Attention

A phenomenon which is vital towards the spread of social media information is "collective attention" [39]. Hence, researchers have been attracted toward understanding how attention to new information propagates among large groups of people. While some studies have shown that dynamics of collective attention of online content is characterized by bursts signifying popularity changes [23, 39], others have demonstrated a natural time scale over which attention fades [50]. A study investigating the emergence of collective attention on Twitter, found that although people's attention is dispersed over a wide variety of concerns, it can concentrate on particular events and shift elsewhere either very rapidly or gradually [42]. Another parallel study focusing on spikes of collective attention in Twitter, analyzed the popularity peaks of Twitter hashtags [23]. They found that the evolution of hashtag popularity over time defined discrete classes of hashtags. Drawing on the progress of these studies, we ask: does the process of evolving collective attention reflect the underlying credibility of a social media story? Unraveling the relation between collective attention rhythms and corresponding credibility level is a complex empirical problem. It requires longitudinal tracking of collective mentions of newsworthy stories in social media along with their in-situ credibility judgments. To that end, CREDBANK provides the most consistently tracked social media information and its associated credibility scores.

2.3 Time Matters

One useful way to understand the interplay between collective attention and information credibility is to examine user activity and information patterns through the lens of time. For decades social scientists have investigated the timing of individual activity to understand the complexity of collective human action. They have reported that timing can range from random [19] to well correlated bursty activity patterns [3]. The bursts in human collective action have not only led to social media reporting of emerging topics, but have also exhibited rich temporal dynamics of social media information spread [31]. For example, information diffusing through micro-blogging platforms like Twitter have demonstrated a short life span [50], with content rising and falling in popularity within hours; whereas, short quoted phrases (known as *memes*) have displayed several days to rise and fade away [25]. On the other hand, general themes (like 'politics', 'economy', 'terrorism') have shown an even larger temporal life span [17, 48]. Social psychologists studying the spread of news and rumor have also noted the importance of temporal patterns in rumor transmission – different types of rumor mongering statements persist over varying temporal spans [5, 43]. However, despite the importance of temporal patterns in information diffusion and rumor transmission, there has been little work in understanding temporal trends in events and its associated credibility assessments. This paper is a step towards unraveling that relation.

3. METHOD

3.1 Data Description

The data investigated in this work was gathered from the CREDBANK corpus [32] which we had built to systematically study social media credibility. The corpus contains 1,377 events as they surfaced on Twitter between October 2014 and February 2015, their corresponding public tweets (a total of 66M messages) and their associated credibility ratings. We built CREDBANK by iteratively tracking millions of public posts streaming through Twitter; computationally detecting the underlying topics (i.e. clusters) of discussion in every block of million tweets; separating event-specific topics from non-event topics by asking independent human raters from Amazon Mechanical Turk (AMT); and then for each of these event topics we gathered credibility ratings on a 5-point Likert scale ranging from 'Certainly Inaccurate' (-2) to 'Certainly Accurate' (+2).

Cred Class	P_{ca} Range	Total Events	Distinct Events
Perfect	$0.9 \leq P_{ca} \leq 1.0$	421	342
High	$0.8 \leq P_{ca} < 0.9$	433	337
Moderate	$0.6 \leq P_{ca} < 0.8$	414	358
Low	$0.0 \leq P_{ca} < 0.6$	109	101

Table 1: Credibility classes and corresponding event counts. "Total Events" column shows event counts from CREDBANK. "Distinct Events" column lists counts from the pruned corpus.

Each event is represented by a combination of top three terms from that event topic. For example, "chelsea", "game", "goal" refers to a football match event at a point in time when Chelsea scored a goal. Thirty independent human raters from AMT judged the accuracy level of an event by browsing through tweets in real-time, where each tweet contained all the top three topical terms. Such a task design ensured that the annotation task closely mimics the way a person would search Twitter to see information related to an event. Moreover, limiting tweets containing all top 3 topical terms ensured a balance between being too generic (by including fewer terms) and too specific (by including more terms), and also provided enough context to a human rater for performing the task.

To guarantee that our collected ratings is at par with expert level judgements, we performed multiple controlled experiments before finalizing the strategy best suited for obtaining quality annotations [33]. Finally, we searched Twitter to collect all tweets specific to the event topic from the past seven days (a time limit enforced by Twitter's search program interface). This did not seem to be a limitation because our experimental setup was tracking recent events. Moreover, research has demonstrated that news worthy topics on twitter have an active period of at most a week [22]. Overall, this iterative framework resulted in a natural experimental setup where the credibility of social media information was being tracked soon after it gained collective attention. Additional details of our data collection process is outlined in the Appendix. A representative sample of events tracked during this three month period along with their credibility ratings is presented in Table 2. By listing the range of diverse events, our aim is to demonstrate the richness of the dataset and hence the generalizability of our results ensuing from this dataset. For example, it contains events ranging from celebrity deaths and terroristic attacks to missing airplanes and soccer matches.

3.2 Pruning Corpus for Sample Independence

During the iterative building of CREDBANK, if an event trended on Twitter for a sufficiently long time period, it is possible that the event is curated multiple times. For example, the event "arsenal", "win", "city" corresponds to the Arsenals winning the football match against Stoke city. People on Twitter had active conversations about the event for several hours, resulting in the event being captured more than once in CREDBANK. However, our statistical analysis (discussed shortly) required sample independence. Occurrence of multiple instances of the same event will likely violate the independence assumption. Hence, we pruned our event sample to keep single distinct instances of each event. By matching the three terms in each event topic, we looked for duplicate event occurrences. Thereafter, if multiple instances of the same event existed, we picked the event which had the earliest curation time. Restricting events by earliest curation times ensured that we retained crowd worker annotations corresponding to the very first time that they performed the annotation task; hence preventing any potential prior knowledge bias. Our pruning step resulted in a dataset of 1,138 events spanning 47,000,127 tweets.

[*] **NOTE**: Credibility assessment is a long running project by the authors of this submission. It started with the creation of CREDBANK corpus [32] and has since been followed with multiple analysis of this corpus [34]. For the ease of the reader, materials and methods have been reproduced in its abbreviated form in the current submission. The analysis, results and the overall contributions of this paper are new.

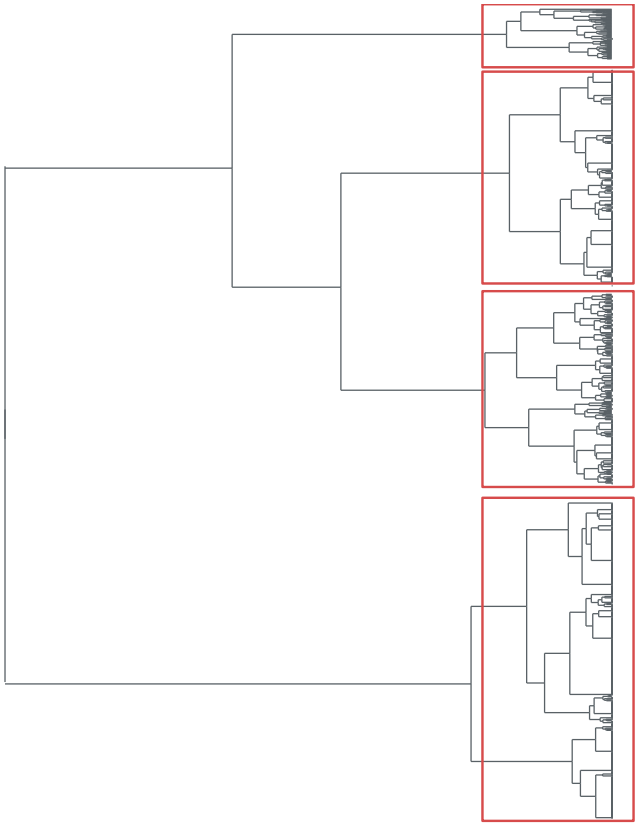

Figure 1: Dendogram from hierarchical clustering of the events from CREDBANK. The boxes show the four clusters.

3.3 Credibility Classification

We measured an event's perceived credibility level based on how many human raters agreed that the event was "Certainly Accurate". More formally, for each event we find the proportion P_{ca} of ratings marked as "Certainly Accurate".

$$P_{ca} = \frac{Count\ "Certainly\ Accurate"\ ratings\ for\ an\ event}{Total\ ratings\ for\ that\ event}$$

To have a reasonable comparison it is impractical to treat P_{ca} as a continuous variable and have a category corresponding to every value of P_{ca}. Hence, we placed P_{ca} into four classes that cover a range of values (see Table 1). The class names are based on the perceived degree of accuracy of the event in that class. For example, events belonging to the "Perfect Credibility" class were rated as "Certainly Accurate" by almost all raters ($0.9 \leq P_{ca} < 1$).

3.4 Validating credibility classification

To ensure that our P_{ca} based credibility classification is a reasonable classification, we compared classes generated by our P_{ca} method against those obtained via data-driven classification.

Event Terms	# Tweets	Start time	End Time	Ratings	P_{ca}
Perfect Credibility: 0.9 ≤ P_{ca} ≤ 1					
george clooney #goldenglobes	10350	2015-01-12 08:50	2015-01-12 18:10	[0 0 1 1 28]	0.93
king mlk martin	88045	2015-01-15 22:00	2015-01-15 22:00	[0 0 0 2 28]	0.93
win pakistan test	5478	2014-10-26 18:10	2014-11-03 21:00	[0 0 0 3 27]	0.90
george arrested zimmerman	45645	2015-01-07 19:40	2015-01-11 00:50	[0 0 0 3 27]	0.90
scott rip sad	26006	2014-12-29 07:50	2015-01-05 18:10	[0 0 0 3 27]	0.90
High Credibility: 0.8 ≤ P_{ca} < 0.9					
beckham odell catches	21848	2014-11-04 04:10	2014-11-04 22:20	[0 0 0 4 26]	0.87
eric garner death	180582	2014-11-26 08:30	2014-12-04 07:10	[1 1 0 2 26]	0.87
windows microsoft holographic	18306	2015-01-21 23:40	2015-01-25 10:00	[0 0 0 4 26]	0.87
kayla mueller isis	65819	2015-02-06 21:10	2015-02-12 00:10	[0 0 0 8 52]	0.87
liverpool arsenal goal	16713	2014-12-14 05:20	2014-12-14 05:20	[0 1 0 4 25]	0.83
Moderate Credibility: 0.6 ≤ P_{ca} < 0.8					
children pakistan #peshawarattack	24239	2014-12-16 12:30	2014-12-17 20:10	[0 1 1 5 23]	0.77
#ericgarner protesters police	12510	2014-12-04 00:50	2014-12-05 10:20	[0 0 2 6 22]	0.73
sydney hostage #sydneysiege	21835	2014-12-15 04:20	2014-12-15 17:20	[0 0 2 6 22]	0.73
bobby shmurda bail	22362	2014-12-17 21:40	2014-12-19 17:30	[0 0 1 7 22]	0.73
#antoniomartin ambulance shot	6330	2014-12-24 11:30	2014-12-24 23:10	[0 0 3 9 18]	0.60
Low Credibility: 0 ≤ P_{ca} < 0.6					
gerrard liverpool steven	204026	2014-12-26 03:40	2015-01-02 20:20	[0 1 3 9 17]	0.57
#chapelhillshooting muslim white	35282	2015-02-11 11:20	2015-02-13 06:20	[2 2 8 16 32]	0.53
paris boko killed	3917	2015-01-07 22:50	2015-01-11 01:50	[0 3 1 11 15]	0.50
ebola #ebola travel	27796	2014-10-09 06:10	2014-10-17 09:10	[2 2 6 10 10]	0.33
killed hostage isis	25925	2015-01-31 20:20	2015-02-08 10:00	[0 8 22 14 16]	0.27

Table 2: Sample events from the CREDANK corpus grouped by their credibility classes. Events are represented with three event terms. Start and end times denote the time period during which tweets were collected using Twitter's search API combined with a search query containing a boolean *AND* of all three event terms. Rating shows the count of Turkers who selected an option from the 5-point Likert scale ranging from -2 ("Certainly Inaccurate") to +2 ("Certainly Accurate").

3.4.1 *Generating data-driven credibility classes*

We used hierarchical agglomerative clustering (HAC) [30] to generate data-driven classes of the credibility rating distributions. HAC is a bottom-up clustering approach which starts with each observation in its own cluster followed by merging pairs of clusters based on a similarity metric. In the absence of a prior hypothesis regarding the number of clusters, HAC is the preferred clustering method. HAC-based clustering approach groups the events based on the shape of their credibility curves on the 5-point Likert scale. Such shape based clustering approach has been used in prior work to cluster based on the shape of popularity peaks [10, 52]. We used the Euclidean distance similarity metric and Ward's fusion strategy for merging [49]. The choice of this strategy minimizes the within-cluster variance thus maximizing within-group similarity [49]. Figure 1 shows the resulting dendogram from hierarchical clustering. The boxes correspond to the credibility groups when the dendogram is cut into four clusters.

3.4.2 *Comparing P_{ca} classes to HAC-based classes*

Is the P_{ca} based credibility classification a close approximation of the HAC based classification? Essentially, we need a metric to compare two clusterings of the same dataset. In other words, we need to measure how often both clustering methods classify the same set of observations as members of the same cluster. We borrow a technique proposed by Tibshirani et al. [46]. Let $P_{clust} = \{x_{1c_1}, x_{2c_1}, x_{3c_2}, \cdots, x_{nc_4}\}$ denote the cluster labels from P_{ca} based classification and $H_{clust} = \{x_{1h_1}, x_{2h_3}, x_{3h_3}, \cdots, x_{nh_4}\}$ the labels from HAC-based classification of the same dataset D of n observations. Here, x_{ic_j} denotes that the i^{th} observation belongs to cluster c_j as per the P_{ca} classification and x_{ih_j} denotes that the i^{th} observation belongs to cluster h_j as per the HAC classification. We see that x_{1c_1} and x_{2c_1} belong to the same cluster. Such pairs are called "co-members". While (x_{1c_1}, x_{2c_1}) are co-members as per P_{ca} classification, (x_{2h_3}, x_{3h_3}) are co-members from HAC classification. For each clustering method, we first compute all pairwise co-memberships. Next, we measure agreement between the clustering methods by computing the Rand similarity coefficient (R) from co-memberships.

$$R = \frac{N_{11} + N_{00}}{N_{11} + N_{10} + N_{01} + N_{00}}$$

N_{11} : number of observation pairs where both are co-members in both clustering methods.
N_{10} : number of observation pairs where the observations are co-members in the first clustering method, but not in the second.
N_{01} : number of observation pairs where the observations are co-members in the second clustering method, but not in the first.
N_{00} : number of observation pairs where neither pair is co-member in either clustering method.

Rand similarity coefficients range between 0 and 1, with 1 corresponding to perfect agreement between the two clustering methods. We obtain a fairly high R of 0.774 denoting high agreement between our P_{ca} based and HAC-based clustering approaches. We favor our proportion-based (P_{ca}) clustering technique over data-driven approaches because the former is much more interpretable and readily generalizable and adaptable to domains other than Twitter on which CREDBANK was constructed.

4. STATISTICAL MEASURES

To understand the relation between collective attention and information credibility, we computed our measures using time-stamped tweets from the CREDBANK corpus, where groups of tweets corresponded to discussions of an event by multiple Twitter users over a certain time span.

4.1 Collective Attention Metrics

Collective attention of such event reportage was measured using two metrics:

1. *Message Volume*: message volume tracks the aggregate number of messages over time
2. *People Volume*: people volume records the aggregate count of unique users paying attention to the story over time.

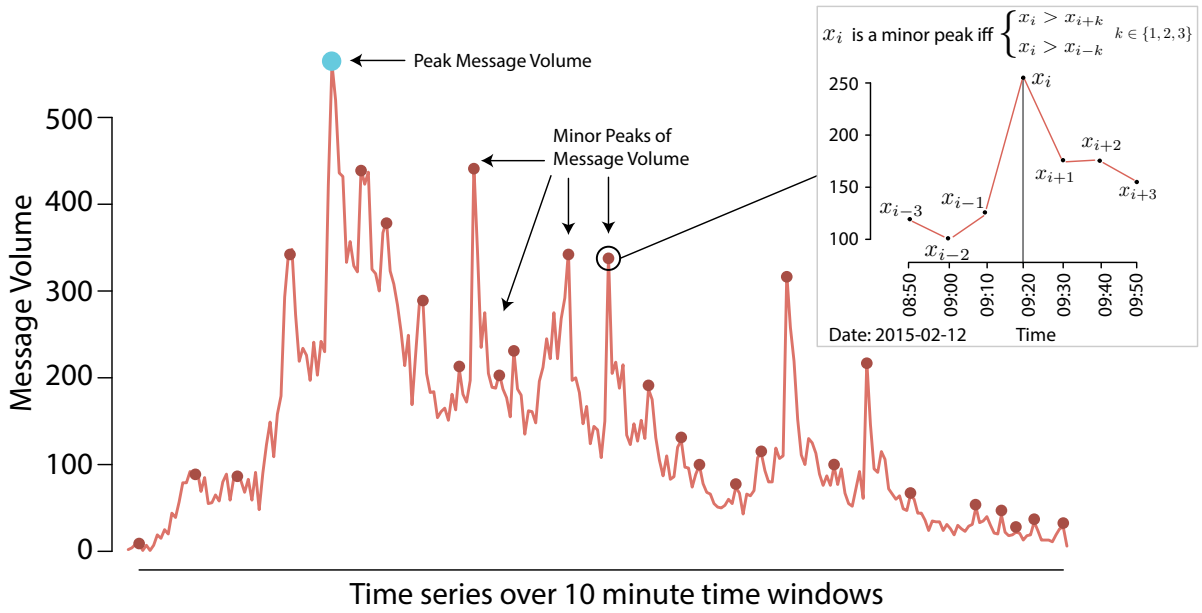

Figure 2: The time series of message volume for a sample event reported on Twitter. The event corresponds to Twitter discussions, where each tweet contained all three terms: "#chapellhillshooting", "muslim" and "white". The ● dot corresponds to the time window having maximum *message volume* while the ● dots correspond to the minor peaks observed in this volume. The inset diagram on the right side zooms in on one of the minor peaks, along with the rule triggering its designation.

Each measure is represented as a time series with message (or unique user) counts aggregated over 10-minute time intervals. Our choice of a 10-minute window is supported by studies showing that Twitter acts as a medium for reporting breaking news and hence is characterized by fast diffusion of information [21, 22]. Thus, tracking collective attention on the order of minutes is a reasonable representation of a rapidly evolving phenomenon. Each event may differ in the temporal dynamics of its collective attention; thus inferences drawn on a small set of events tracked for a few days may be confounded by temporal traits peculiar to certain news stories. However, by tracking news stories over several months and averaging over hundreds of such collective attention rhythms, our results represent the most consistent relations between the dynamics of collective attention and perceptions of information credibility. Our rationale for using *people volume*, in addition to *message volume*, as a collective attention metric is to ensure that the collective attention measured is not confounded by superfluous posting activity from potential Twitter bots, automated programs posing as human beings [8]. Since *people volume* corresponds to the unique number of individuals paying attention to the event over time, it aims to counteract any extreme posting activity by such bots.

As an example illustrating our data and methods, Figure 2 shows aggregate message volume for an event reported on Twitter where every message contained the terms "#chapellhillshooting", "muslim" and "white". On February 10, 2015, three Muslim students in Chapel Hill, North Carolina were shot to death by a white neighbor and speculations concerning the motives of the shooter surrounded the event [14]. While authorities suggested the motive to be an ongoing dispute between neighbors over a parking space, many social media users suggested a hate crime as the motive. Twitter messages concerning this specific topic on February 12 blamed the media for ignoring the coverage of an event involving Muslim killings and suggested the shooting was an act of terrorism and so a hate crime. It was not until February 13 that authorities opened an investigation to determine if the shooting was in fact a hate crime. Credibility rating distributions showed that less than 50% of raters agreed that the social media reportage of the event was "Certainly Accurate", thereby questioning the alleged terror claims underlying the act.

4.2 Temporal Measures of Collective Attention

To quantify the importance of the time when collective attention maximized, we first computed the strict global maximum in the time series [45]. We call this the *peak attention*. This is the ratio of messages (or unique people) within the peak time window to the total cumulative volume of messages (or unique people) over the entire event time series:

$$\textit{Peak Attention} = \frac{max(x_1, \cdots, x_n)}{\sum_{i=1}^{n} x_i} \quad (1)$$

where x_i is the count of messages (or unique people) in time window i in an event time series $x_1, x_2, \cdots, x_n$. Our choice of the above measure is based on the success of prior studies using peak fraction based metrics to successfully characterize herding behavior over time [10, 52]. To illustrate how peak attention measure can characterize variations in time series, consider the example of an event reportage marked by a sudden spike in collective attention followed by a subsequent drop. The lack of precursory growth suggests that most of the attention was concentrated on the peak, thereby resulting in high *peak attention* (Figure 4c and 4d). Whereas, an event with steady growth in collective attention, followed by a gradual decay would imply a relatively smaller fraction of attention in the peak, thus leading to lower value of *peak attenion* (Figure 4b).

While *peak attention* captures the importance of the maximal burst in collective attention, it does not take into account the presence or absence of spikes in the precursory growth and in the subsequent decay following the burst. Hence, we define a measure to quantify the spikiness in collective attention. We detect all strict local maxima [45] in each of the event time series. A strict local maxima corresponds to an instance in the time series when the volume of messages (or unique people) is larger than the volume in the neighboring time windows. We define this neighborhood as three

time windows on either side of the local maxima and call these local maxima *minor peaks*. Thus, the attention in a minor peak is higher than the attention 30 minutes (i.e., three time windows times 10-minute window size) before and after the occurrence of a peak.

$$x_i \text{ is a minor peak iff } \begin{cases} x_i > x_{i+k} \\ x_i > x_{i-k} \end{cases}, k \in \{1, 2, 3\} \quad (2)$$

We then define minor peak attention as the ratio of messages (or unique people) in the local maxima relative to the total cumulative volume of messages (or unique people) over the entire event time series. Formally, if $\mathcal{M}$ is the set of all minor peak indices in an event's message (or unique people) time series, then minor peak attention is defined as follows:

$$\textit{Minor Peak Attention} = \frac{\sum_{j \in \mathcal{M}} x_j}{\sum_{i=1}^{n} x_i} \quad (3)$$

The inset diagram in Figure 2 shows a local maximum. While the peak attention captures the maximum momentary interest that an event acquires during its lifetime on Twitter, the points representing minor peak attention reflect renewed and ongoing recurrences of momentary interest. Additionally, both these measures have two important properties: both are invariant with respect to scaling and shifting [52]. First, since both measures are proportions based on cumulative collective attention, they are invariant to the overall volume of attention. Hence, two event time series having similar peaky shapes but different total attention volumes would be treated similarly. Secondly, both measures are computed independent of the maxima position on the time axis. Thus, if two event time series peaks occur at different times but possess a similar peaky structure, the measures will be invariant to the translations on the time axis. Hence, both these measures—despite being simple representations of temporal dynamics—are useful in interpreting the relationship between collective attention rhythms and event credibility across a range of different events exhibiting high variability in overall popularity and time of popularity.

4.3 Statistical Analysis and Results

We tested the differences in collective attention measures across the credibility classes using the Wilcoxon Rank Sum or Mann-Whitney U test. For each temporal measure (peak and minor peak attention) and for each collective attention metric (message and people volume), we performed pairwise Wilcoxon Rank Sum tests, followed by Bonferroni corrections [12] to control for potential inflation of the family-wise error rate by multiple test comparisons. We found that, for both *message volume* and *people volume*, differences in the minor peak fraction are statistically significant ($p < 0.00833$ after Bonferroni corrections and using Wilcoxon Rank Sum tests). As shown in Figure 3, median *minor peak attention* decreases as credibility level increases from "Low" to "Perfect". We also found a significant moderate degree of negative correlation between P_{ca} and minor peak fraction for both *message volume* ($r = -0.33$) and *people volume* ($r = -0.33$). These results suggest that an event attracting renewed interest is associated with lower perceived credibility. On the other hand, *peak attention* of messages was only statistically different between "Perfect" and "Moderate" credibility classes. Peak attention for *people volume* could only provide coarse-grained information separating "Low" and "Perfect", and "Moderate" and "Perfect" credibility classes. These results indicate that *peak attention* is not a useful signal for event credibility. To ensure that these ratio-based, collective attention measures described above are not sensitive to event duration, which can affect the denominator (cumulative volume), we performed pairwise Wilcoxon Rank Sum test comparisons of event duration across the credibility classes. We found no significant difference, indicating that event duration does not skew the collective attention metrics for a particular credibility class. Moreover, to ensure that our collective attention metrics are independent observations over time–a criteria necessary for the validity of our statistical analysis–we performed Ljung-Box Q (LBQ) tests [27]. We were able to reject null hypothesis for each of our LBQ tests; thus confirming that our collective attention measures for both message and people volume are independent over time.

	P	H	M	L
P		0.00012**	5.4e-12****	4.6e-11****
H	0.002639*		0.001703*	3.6e-06****
M	1.0e-08****	0.004221**		0.00547*
L	8.4e-11****	1.5e-06****	0.001981**	

(a) Minor peak attention pairwise statistical differences

	P	H	M	L
P		ns	0.00322*	ns
H	ns		ns	ns
M	0.00093***	ns		ns
L	0.00045***	ns	ns	

(b) Peak attention pairwise statistical differences

Table 3: Pairwise statistical significance after Wilcoxon Rank Sum tests. P, H, M, L correspond to Perfect, High, Moderate and Low credibility classes. The top half of the diagonal corresponds to *message volume*, while bottom half shows pairwise differences in *people volume*. ns stands for non-significance.

5. DISCUSSION

By investigating the most comprehensive large-scale longitudinal credibility corpus constructed to date, we were able to test the relationship between an event's perceived credibility level and the temporal dynamics of its collective attention. According to our findings, moments of renewed collective attention are associated with event reportage marked by decreased levels of perceived credibility. Do frequent peaks in collective attention lead to lower perceived credibility? Or do reduced levels of credibility spark the continued interest in the event? Our current study cannot establish the causal direction of this relation. However, we are able to establish that the persistence of collective attention peaks is a reliable temporal signature for an event's perceived credibility level. Moreover, an advantage of viewing these phenomena through the lens of a fundamental property of human activity, such as time, is that the resultant findings are likely to hold irrespective of the platform (e.g., Twitter) hosting the collective human attention directed toward real-world events.

We remark that by using a simple proportion based classification technique we identified robust and scalable credibility classes; hence it is also potentially applicable to other online settings where user's collective attention drives popularity of content. Moreover, by using simple and interpretable parameters computed on times series of minute-wise user and message attention, we revealed vital temporal indicators associated with information credibility. Contrary to other sophisticated methods which require the estimation of power-law exponents for unraveling collective attention dynamics, or the calculation of costly correlations between activity time series, the parameters employed here can easily be computed in a scalable way. Although devoid of any predictive power, these measures can support the discovery of collective attention patterns in large-scale records of human activity.

On the basis of these results, we envision that organizations struggling to handle the propagation of online misinformation [28] can

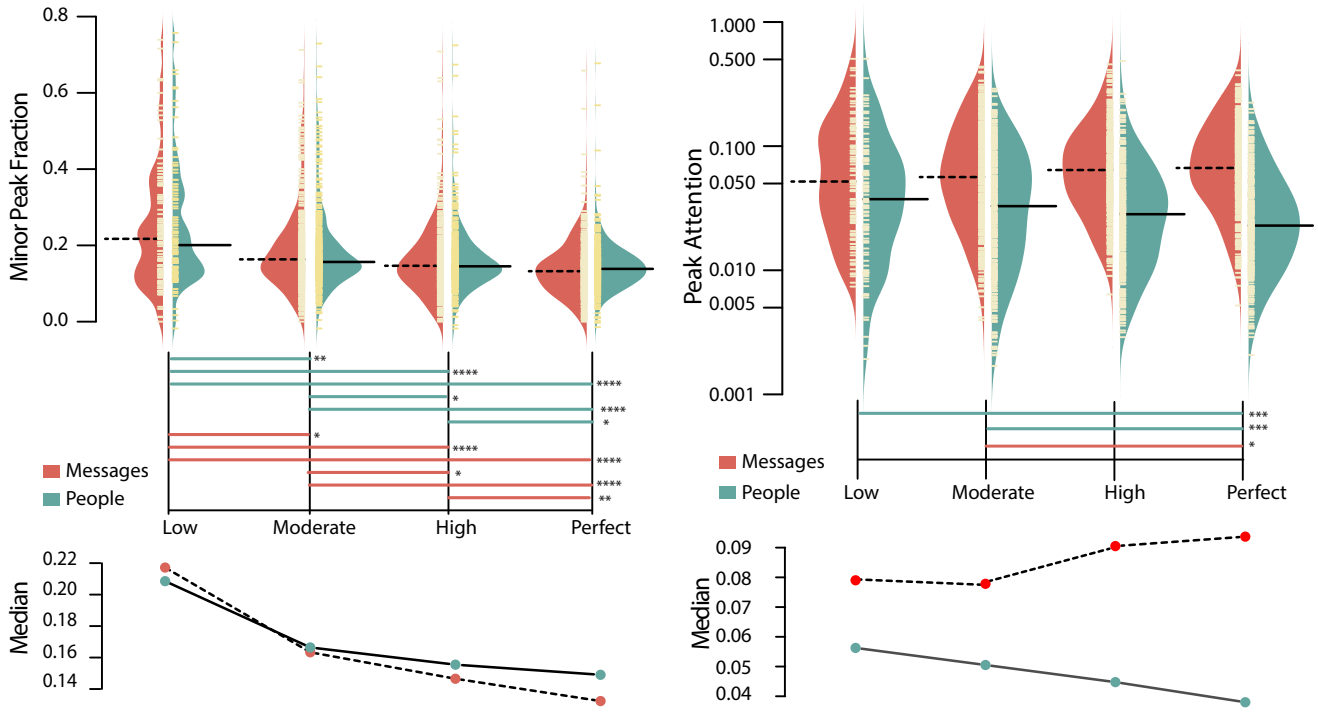

(a) (b)

Figure 3: Collective attention shown as a beanplot distribution. The shape of each half of the asymmetric bean represents the Gaussian density estimation of the distribution. The lines (in yellow) are the actual data points; the dotted long bean line is the median corresponding to the message volume, the solid line shows the median for people volume. The * denotes pairwise significant differences between cluster medians after correcting for familywise error-rate. (a). Proportion of minor peak fractions are statistically different across all credibility class pairs for both message and unique people volume. (b). Peak attention is significantly different across "Low" and "Perfect", and "Moderate" and "Perfect" credibility classes for unique people volume, and "Perfect" and "Moderate" classes for message volume. The line charts at the bottom panel show the median trends across the credibility classes.

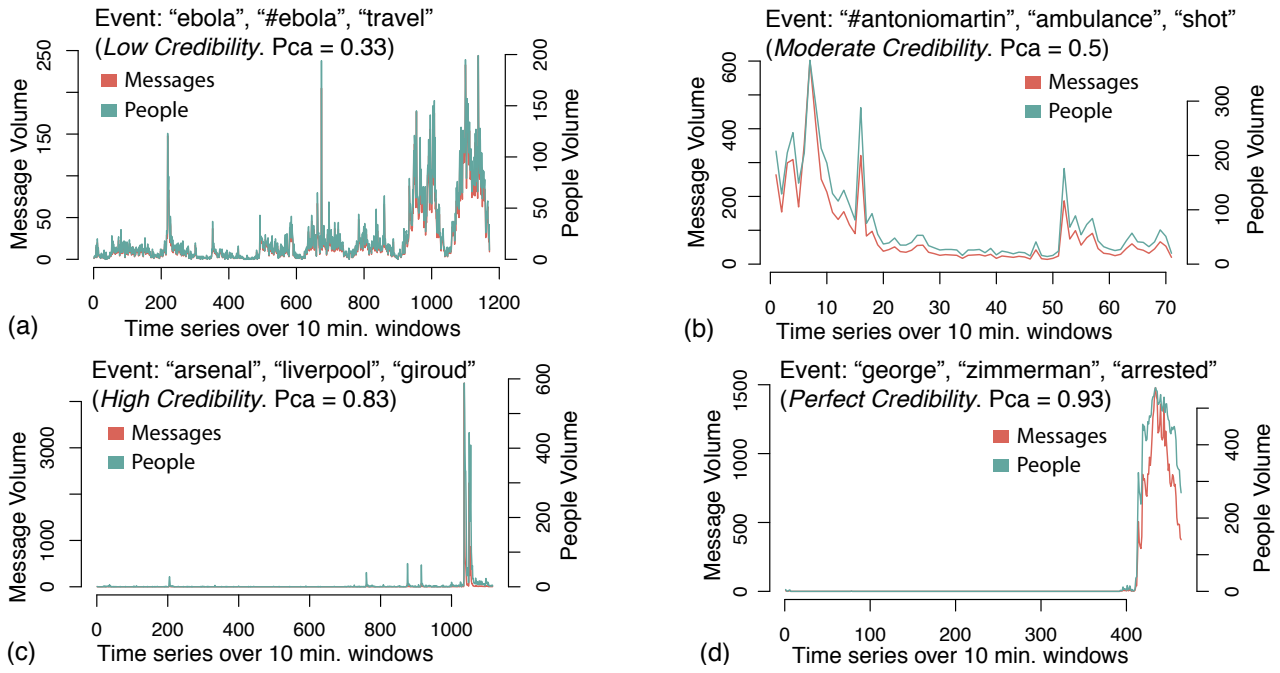

Figure 4: Time series of collective attention metrics (message volume and unique people volume) for example events in each credibility class. The examples show representative behavior of collective attention metrics in each credibility class. While events in all four classes are marked by peak attention with respect to both message and people volume, events in the low and moderate credibility classes exhibit multiple minor peaks, signifying that persistent attention is characteristic of lower credible social media events.

harness the temporality of collective attention to predict the level of credibility. We may be able to subsequently design interventions aimed at controlling the spread of false information or cautioning social media users to be skeptical about an evolving topic's veracity, ultimately raising an individual's capacity to assess credibility of information. Imagine a news reporting tool which shows social media discussions highlighting areas which witnessed multiple minor peaks of human activity, or think of a fact-checking system that compares temporal regions of high minor peak attentions to those with fewer attention peaks, or consider temporal tagging of scientific discourse or medical records emphasizing areas that garnered intermittent temporal popularity. We foresee that our findings can lead to a new class of such temporally aware systems which underscore degrees of information uncertainty based on temporal signals of collective attention. Finally, our study has practical implications in the field of computational social science where inferences about human social behavior are based on reports of online interactions [9] and trusting the credibility of those reports is crucial for any downstream analysis. For example, imagine a health researcher investigating the spread of Ebola via social media reports or a financial trader gauging market volatility based on breaking news and citizen reports on social media; veracity of those reports will affect the subsequent inferences.

6. CONCLUSIONS

To study the dynamics of collective attention and its relation to information credibility in a natural setting, we analyzed the temporal patterns of 47M Twitter messages spread across 1,138 social media events along with their in-situ credibility ratings. We do so by multiple statistical comparison tests over parameters computed on the time series of collective attention of messages and distinct users. Although simple, this approach provides fundamental insights about collective attention and information credibility that would otherwise be missed by more complicated predictive analysis methods.

7. REFERENCES

[1] A. Arif, K. Shanahan, F. Chou, Y. Dosouto, K. Starbird, and E. S. Spiro. How information snowballs: Exploring the role of exposure in online rumor propagation. In *Proc. CSCW*, pages 465–476, 2016.

[2] S. Asur, B. A. Huberman, G. Szabo, and C. Wang. Trends in social media: Persistence and decay. *Available at SSRN Scholarly Paper ID 1755748, Social Science Research Network.*, 2011.

[3] A.-L. Barabasi. The origin of bursts and heavy tails in human dynamics. *Nature*, 435(7039):207–211, 2005.

[4] M. Bass. Frank. *A New Product Growth for Model Consumer Durables*, 50:1825–1832, 1969.

[5] P. Bordia and R. L. Rosnow. Rumor rest stops on the information highway transmission patterns in a computer-mediated rumor chain. *Human Communication Research*, 25(2):163–179, 1998.

[6] C. Castillo, M. Mendoza, and B. Poblete. Information credibility on twitter. In *Proc. WWW*, 2011.

[7] A. Caumont. *12 trends shaping digital news* (Pew Research Center). http://www.pewresearch.org/fact-tank/2013/10/16/12-trends-shaping-digital-news, 2013.

[8] Z. Chu, S. Gianvecchio, H. Wang, and S. Jajodia. Who is tweeting on twitter: human, bot, or cyborg? In *Proc. ACSAC*, pages 21–30. ACM, 2010.

[9] S. Counts, M. De Choudhury, J. Diesner, E. Gilbert, M. Gonzalez, B. Keegan, M. Naaman, and H. Wallach. Computational social science: Cscw in the social media era. In *Proc. CSCW Companion publication*, pages 105–108. ACM, 2014.

[10] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *PNAS*, 105(41):15649–15653, 2008.

[11] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi. The spreading of misinformation online. *PNAS*, 113(3):554–559, 2016.

[12] O. J. Dunn. Estimation of the medians for dependent variables. *The Annals of Mathematical Statistics*, pages 192–197, 1959.

[13] A. Friggeri, L. Adamic, D. Eckles, and J. Cheng. Rumor cascades. In *Proc. ICWSM*, 2014.

[14] S. Frizell. *3 Muslim Students Murdered in North Carolina*. time.com/3704759/muslim-students-murdered-chapel-hill, 2015.

[15] S. Goel, J. M. Hofman, S. Lahaie, D. M. Pennock, and D. J. Watts. Predicting consumer behavior with web search. *PNAS*, 107(41):17486–17490, 2010.

[16] S. A. Golder and M. W. Macy. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051):1878–1881, 2011.

[17] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through blogspace. In *Proc. WWW*, pages 491–501. ACM, 2004.

[18] A. Gupta, P. Kumaraguru, C. Castillo, and P. Meier. Tweetcred: Real-time credibility assessment of content on twitter. In *International Conference on Social Informatics*, pages 228–243. Springer, 2014.

[19] F. A. Haight and F. A. Haight. Handbook of the poisson distribution. 1967.

[20] L. Howell. Digital wildfires in a hyperconnected world. *World Economic Forum*, 2013.

[21] M. Hu, S. Liu, F. Wei, Y. Wu, J. Stasko, and K.-L. Ma. Breaking news on twitter. In *Proc. CHI*, pages 2751–2754. ACM, 2012.

[22] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proc. WWW*, pages 591–600. ACM, 2010.

[23] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto. Dynamical classes of collective attention in twitter. In *Proc. WWW*, pages 251–260. ACM, 2012.

[24] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.

[25] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proc. SIGKDD*, pages 497–506. ACM, 2009.

[26] F. Liu, A. Burton-Jones, and D. Xu. Rumors on social media in disasters: Extending transmission to retransmission. In *Proc. PACIS*, 2014.

[27] G. M. Ljung and G. E. Box. On a measure of lack of fit in time series models. *Biometrika*, 65(2):297–303, 1978.

[28] V. Luckerson. *Fear, Misinformation, and Social Media Complicate Ebola Fight.* Time Inc. time.com/3479254/ebola-social-media/, 2014.

[29] J. Maddock, K. Starbird, H. J. Al-Hassani, D. E. Sandoval, M. Orand, and R. M. Mason. Characterizing online rumoring

behavior using multi-dimensional signatures. In *Proc. CSCW*, pages 228–241. ACM, 2015.

[30] O. Maimon and L. Rokach. *Data mining and knowledge discovery handbook*, volume 2. Springer, 2005.

[31] M. Mathioudakis and N. Koudas. Twittermonitor: trend detection over the twitter stream. In *Proc. SIGMOD*, pages 1155–1158. ACM, 2010.

[32] T. Mitra and E. Gilbert. Credbank: A large-scale social media corpus with associated credibility annotations. In *Proc. ICWSM*, 2015.

[33] T. Mitra, C. J. Hutto, and E. Gilbert. Comparing person-and process-centric strategies for obtaining quality data on amazon mechanical turk. In *Proc. CHI*, pages 1345–1354. ACM, 2015.

[34] T. Mitra, G. Wright, and E. Gilbert. A parsimonious language model of social media credibility across disparate events. In *Proc. CSCW*, 2017.

[35] M. R. Morris, S. Counts, A. Roseway, A. Hoff, and J. Schwarz. Tweeting is believing?: understanding microblog credibility perceptions. In *Proc. CSCW*, pages 441–450. ACM, 2012.

[36] M. Naaman, J. Boase, and C.-H. Lai. Is it really about me?: message content in social awareness streams. In *Proc. CSCW*, pages 189–192. ACM, 2010.

[37] B. O'Connor, M. Krieger, and D. Ahn. Tweetmotif: Exploratory search and topic summarization for twitter. In *Proc. ICWSM*, 2010.

[38] V. Qazvinian, E. Rosengren, D. R. Radev, and Q. Mei. Rumor has it: Identifying misinformation in microblogs. In *Proc. EMNLP*, 2011.

[39] J. Ratkiewicz, S. Fortunato, A. Flammini, F. Menczer, and A. Vespignani. Characterizing and modeling the dynamics of online popularity. *Physical review letters*, 105(15):158701, 2010.

[40] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc. WWW*, pages 695–704. ACM, 2011.

[41] M. J. Salganik, P. S. Dodds, and D. J. Watts. Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854–856, 2006.

[42] K. Sasahara, Y. Hirata, M. Toyoda, M. Kitsuregawa, and K. Aihara. Quantifying collective attention from tweet stream. *PloS one*, 8(4):e61823, 2013.

[43] T. Shibutani. *Improvised news: A sociological study of rumor*. Ardent Media, 1966.

[44] S. Sinha, A. Chatterjee, A. Chakraborti, and B. K. Chakrabarti. *Econophysics: an introduction*. John Wiley & Sons, 2010.

[45] J. Stewart. *Calculus: Early transcendentals*. Cengage Learning, 2015.

[46] R. Tibshirani and G. Walther. Cluster validation by prediction strength. *Journal of Computational and Graphical Statistics*, 14(3):511–528, 2005.

[47] Z. Tufekci. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *In Proc. ICWSM*, 2014.

[48] X. Wang, C. Zhai, X. Hu, and R. Sproat. Mining correlated bursty topic patterns from coordinated text streams. In *Proc. SIGKDD*, pages 784–793. ACM, 2007.

[49] J. H. Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244, 1963.

[50] F. Wu and B. A. Huberman. Novelty and collective attention. *PNAS*, 104(45):17599–17601, 2007.

[51] Y. Wu, C. Zhou, J. Xiao, J. Kurths, and H. J. Schellnhuber. Evidence for a bimodal distribution in human communication. *PNAS*, 107(44):18803–18808, 2010.

[52] J. Yang and J. Leskovec. Patterns of temporal variation in online media. In *Proc. WWW*, pages 177–186. ACM, 2011.

[53] L. Zeng, K. Starbird, and E. S. Spiro. # unconfirmed: Classifying rumor stance in crisis-related social media messages. In *Proc. ICWSM*, 2016.

APPENDIX

A. CREDBANK CONSTRUCTION STEPS

CREDBANK was constructed by following a sequence of phases:

1. Streaming Tweets and Preprocessing: Twitter's Streaming API was used to iteratively collect a continuous 1% sample of all global tweets, filtered to contain only English tweets, followed by spam removal, tokenization using a Twitter specific tokenizer [37] and a sophisticated multi-stage stop word removal step.

2. Detecting Event Candidates: After carefully considering various approaches for event detection from social media streams, we opted for topic models, since topic models can learn term co-occurences and unlike keyword based techniques do not make a-priori assumptions about what constitutes an event.

3. Event annotation: To eliminate detection of potential false positives using a purely computational event detection approach, candidate events from the previous step were sent to ten independent human raters from Amazon Mechanical Turk (AMT) for judging whether a topic relates to a real-world news event. The majority agreement was selected as the final annotation.

4. Credibility Assessment: This phase had three primary steps:

Determining the credibility scale: Informed by work done by the linguistic community on 'Event Factuality', the credibility scale was designed as an interaction between two dimensions: Polarity, which differentiates among 'Accurate', 'Inaccurate', and 'Uncertain', and Degree of certainty which distinguishes among 'Certainly','Probably' and 'Uncertain', leading to a 5-point Likert scale annotation scheme.

Determining number of independent Turk ratings for high quality annotation: In this step, we piloted the CREDBANK system for 5 days collecting and annotating 50 events by both Turkers and expert annotators (university research librarians). The pilot study was followed by computing correlation statistics between Turker mean responses and expert mean responses while varying the count of independent Turker ratings per event. The correlation maximized at 30 Turker ratings leading to the decision of collecting 30 annotations per event.

Credibility assessment task: The credibility assessment task framework was designed to ensure that the collected credibility ratings is of high quality. Multiple controlled experiments were performed before finalizing the strategy best suited for obtaining quality annotations [33]. Turkers were first selectively screened and trained via a qualification test. Screened workers were then directed to a task interface and asked to categorize an event's credibility after reading through a stream of real-time tweets related to an event topic. They were instructed to either be knowledgeable on the event topic or search online before making their credibility judgments.

5. Collecting Event Streams: The final phase used Twitter's search API to collect all tweets specific to the event topic.